Image Engine: An Object-Oriented Multimedia Database for Storing, Retrieving and Sharing Medical Images and Text

Henry J. Lowe, M.D.

Department of Medicine, Section of Medical Informatics
University of Pittsburgh, Pittsburgh PA

ABSTRACT

This paper describes Image Engine, an object-oriented , microcomputer-based, multimedia database designed to facilitate the storage and retrieval of digitized biomedical still images, video, and text using inexpensive desktop computers. The current prototype runs on Apple Macintosh computers and allows network database access via peer to peer file sharing protocols. Image Engine supports both free text and controlled vocabulary indexing of multimedia objects. The latter is implemented using the TView thesaurus model developed by the author. The current prototype of Image Engine uses the National Library of Medicine's Medical Subject Headings (MeSH) vocabulary (with UMLS Meta-1 extensions) as its indexing thesaurus.

INTRODUCTION

Medicine generates a large volume of text-based data in the form of clinical records, scientific papers and textbooks. The response to this "information explosion" [1] has centered on computer and information technology solutions such as bibliographic databases and clinical information systems. Many of these systems are entirely or largely text-oriented. However text is only one medium through which biomedical information is communicated. In the clinical domain many text-based documents are written descriptions and interpretations of non-textual data such as radiology, pathology, dermatology, ophthalmology and endoscopy images. In addition, the rapid growth of diagnostic imaging technologies [2] over the past two decades has dramatically increased the amount of non-textual data generated in clinical medicine. These images are often seen only by technicians and reporting consultants before being archived in "hard-copy" format on x-ray film, magnetic media or videotape. Improving the clinician's access to diagnostic images may enhance decision making [3]. In addition, providing easy electronic access to a database of selected teaching images would be useful to medical students and educators [4].

While a number of videodisk-based image collections are widely used as educational resources [5, 6] the analog nature of this medium makes it difficult and expensive to access videodisk images over a network. In addition creation and updating of videodisk-based image collections is expensive. Digital imaging technology offers many advantages such as network access, image compression, image processing and image display on inexpensive microcomputers without the need for additional, expensive hardware such as frame grabbers and external video monitors.

While Picture Archiving and Communication Systems (PACS) [7, 8] have been an active area of research and development for almost twenty years much of that work has been within the domain of radiology and has focused on expensive institutional systems. Recent advances in digital image compression, real time digital video, networking, storage media and personal computer technology coupled with the decreasing cost of these ubiquitous systems may well provide the means for affordable universal access to multimedia materials not only in radiology but also in many other areas of medicine.

This paper describes Image Engine, a prototype microcomputer-based system for the storage, retrieval, sharing and manipulation of medical multimedia materials. Image Engine is not intended to compete with institutional PACS. Instead, it was designed to provide an inexpensive solution for small groups and departments to share multimedia materials in support of medical education, research, clinical practice and patient teaching. Given the decreasing cost and likely widespread availability of integrated microcomputer-based image processing technology one may anticipate the increasing capture and use of digital images in these areas. In addition the emergence of "personal multimedia libraries" is likely as individual physicians explore this increasingly inexpensive and sophisticated technology.

0195-4210/92/$5.00 © 1994 AMIA, Inc.

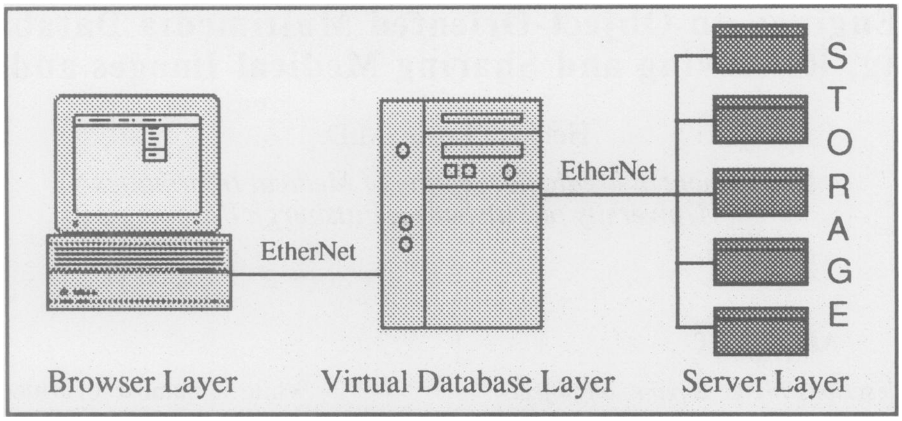

Figure 1: Image Engine Database Architecture

BACKGROUND

Image Engine was designed as a low cost system to permit the storage, retrieval, sharing and manipulation of digitized still images, digital video, and text using inexpensive desktop computer technology. The current prototype was initially developed to support the creation of a digital multimedia library for use by faculty and students in the University of Pittsburgh's new problem-based medical curriculum. Faculty will be invited to submit relevant still images or video and to supply a textual annotation for each item. The materials are digitized, compressed and entered into the Image Engine database along with the textual annotation. The contributor's name and department are also stored with each image. Faculty and students can search for, retrieve, view and copy images from the database using the Image Engine Browser, a stand alone application. It is hoped that this model, which is still in the development phase, will facilitate the creation of multimedia teaching software by faculty and students and also support the use of multimedia resources in the small group sessions which are part of the new medical curriculum.

SYSTEM ARCHITECTURE

The Image Engine multimedia database system consists of 3 layers (Figure 1): the File Server layer, the Object Database layer and the Browser layer.

1. The File Server layer consists of one or more physical storage devices (mainly gigabyte hard disks but CD-ROM, Bernoulli and magneto-optical disks are also supported) located on the network. The server stores still images (usually JPEG compressed PICT files), digitized video (compressed Quicktime files) and text (standard ASCII text files). These files are stored in standard operating system format.

2. The Object Database is an object-oriented database consisting of multimedia objects representing image, video and text objects stored on the File Server. Objects in the database contain slots which store information such as the object's type, name, size, description, thumbnail (a 100 x 100 bit scaled version of the actual image), contributor and up to 32K of associated text. In addition each object contains a slot for a list of indexing terms derived from MeSH. Multimedia objects contain methods that respond to messages sent by the browser. Examples include methods to display the object, assemble and display a database record and copy the object's server file to a user-specified location.

The Object Database layer provides the user with a virtual database in which the multimedia files that reside on the File Server's storage volumes are represented as objects in an apparently local database. This model has a number of advantages. It provides a high degree of data-abstraction which will facilitate any future cross-platform access. It also permits the File Server to actually encompass multiple physical storage devices in different locations on the network.

The user interacts with the Object Database via the Image Engine Browser application. The Object Database in turn handles all interactions with the File Server thus insulating the user from the underlying operating system and networking protocols.

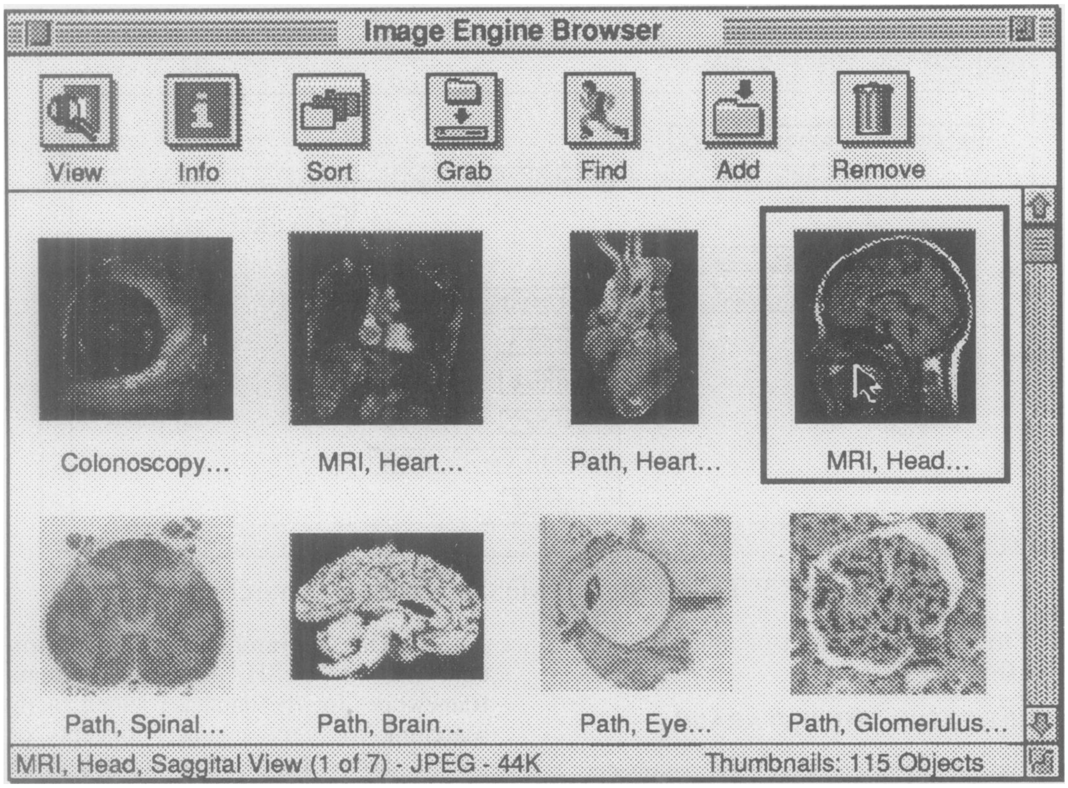

Figure 2. Image Engine Browser

3. The Image Engine Browser (Figure 2) is an application that allows users to interact with the Object Database. Using the browser one can search the database for objects by name, type, size, contributor, description, annotation and indexing terms. Substring and global searches are supported. A search usually results in the retrieved subset of multimedia objects being displayed in either a thumbnail or text list view.

Using the browser, objects may be selected and requested to display themselves. Multimedia objects respond appropriately by displaying the contents of their server file in a separate window. In the case of still images the picture can be scaled, copied to the clipboard, copied to a file on the user's machine or opened by a user-selected application such as an image processor. Digitized video clips may be played on the computer's screen (supports rewind, fast forward, pause, frame forward, frame back). Individual video frames may be converted to a still image which can be scaled, copied to the clipboard, copied to a file on the user's machine or opened by a user-selected application such as an image processor.

The user can also view a database record for each object. The record displays the objects name, type, size, description, text annotation, thumbnail image and indexing terms. Multiple objects and their records may be displayed simultaneously. The browser also supports a Carousel object which represents a linked list of multimedia objects selected by the user. Carousels provide a convenient way for a user to store, retrieve and share groups of objects. A sparse Carousel file containing pointers into the Object Database can then be stored on the users machine and loaded into the browser as needed at a later time. Carousel objects can apply operations to their constituent multimedia objects. The browser currently supports only one such operation which creates a slide show of a Carousel's objects.

MANAGING THE DATABASE

A special administrator's version of the Image Engine Browser supports a number of functions which allow adding, removing, compressing and editing multimedia objects. Objects can be added to the database individually or in groups. The browser can be instructed to scan storage volumes or folders in the Server layer for appropriate object files. The browser will then automatically build multimedia objects (creating thumbnail images if appropriate) and add them to the Object Database. Objects can be removed from the database singly or in groups. The administrator can use the browser's search

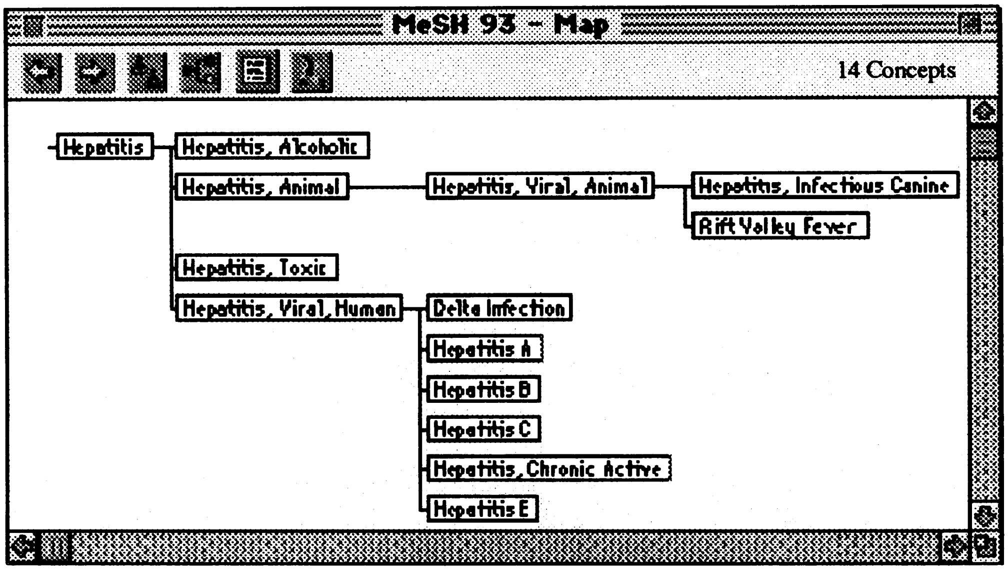

Figure 3: TView Map of MeSH Hepatitis Hierarchy

function to identify a subset of objects to be removed from the database.

INDEXING THE DATABASE

Image Engine can search for multimedia objects using any object data, including the objects annotation. An annotation is usually created by the objects contributor and may be up to 32000 characters is size. It may contain anything from a brief one paragraph description of the object to a multipage document. Therefore objects can be retrieved using a free text search for one or more words in the annotation. However, free text searching may introduce a number of difficulties. The individual writing the annotation may describe a concept in language that differs from the language used by a searcher hoping to retrieve objects related to that concept. Common examples of this phenomenon are the uses of synonymous words such as kidney/renal, liver/hepatic and cancer/neoplasm. Furthermore the same concept may be represented with synonymous terms (cardiac ultrasound / echocardiography), eponymous terms (Wilson's Disease / Hepatolenticular Degeneration) and acronyms (AIDS / Acquired Immunodeficiency Syndrome). For these and other reasons free text searches may result in poor search recall [9]

$$\text{Recall} = \frac{\text{No. of Relevant Items Retrieved}}{\text{No. of Relevant Items in the Database}}$$

Use of a controlled indexing vocabulary for indexing and searching can address these problems and is standard practice in many bibliographic databases. For example, the MEDLINE database is indexed with MeSH.

The hierarchical organization of concepts in MeSH provides a semantic framework for the database which facilitates searching for very specific concepts or for general classes of related concepts (figure 3). For example one can search for a very specific type of viral hepatitis (e.g. Delta Infection), any viral hepatitis or any hepatitis (alcoholic, toxic, viral etc.).

For these reasons it was decided to index each multimedia object with one or more MeSH terms. Using TView, an object-oriented thesaurus browser, the multimedia object's contributor or database administrator can rapidly find relevant MeSH terms for use as indexing terms. TView can also be used to find MeSH terms for inclusion in Image Engine Browser queries.

TVIEW

TView is a Macintosh-based object-oriented thesaurus browser developed by the author. TView builds upon the MicroMeSH [10] and AURIC browsers [11]. Using TView one can navigate the hierarchical and contextual relationships between thesaurus concepts. In addition TView has a search function for finding thesaurus concepts. The TView MeSH thesaurus was build from the Unified Medical Language Systems (UMLS) Metathesaurus. In addition to MeSH data it also contains additional UMLS information including definitions, semantic types and synonyms. TView can be used from within the Image Engine browser and data can be passed between the two applications using the clipboard. TView can be used both for indexing multimedia objects and for finding MeSH terms with which to search the multimedia database.

TView can also be used to retrieve related multimedia objects. For example, if a retrieved object is indexed with the MeSH term DELTA INFECTION then TView will report that this is a type of VIRAL HEPATITIS and display other disease in that group (figure 3). One or more of these MeSH terms can be copied into the Image Engine query builder and a search for related objects executed.

LIMITATIONS

Image Engine's currently runs only on the Macintosh family of computers. However, the separation of the server and virtual database layers reflects a desire to ensure that the multimedia files be in a standard format that can be accessed from other platforms such as Microsoft Windows and UNIX. Macintosh PICT files can be created and opened by many non-Macintosh applications. Macintosh Quicktime movie files can be played using Quicktime for Windows (utilities also exist to convert files between Microsoft Video for Windows and Windows/Macintosh Quicktime formats) and of course ASCII text is the lingua franca for text data exchange. While no cross-platform development is currently planned such development will be much easier given the object-oriented design, data abstraction and standard data formats used by the Image Engine database model.

POTENTIAL USES

Image Engine has a number of potential applications in medicine. It could be used to create an institutional or departmental multimedia library in which users both contribute and share images and text. Many departments and individuals already have collections of 35mm slides, radiographs, videotapes etc. that if digitized, annotated and shared electronically would create a powerful educational resource.

Image Engine can be used by clinical departments to maintain a local database of patient data, still images and video clips of diagnostic materials such as gastroscopy, colonoscopy, fundoscopy, bronchoscopy, pathology or dermatology images. For example a Gastroenterology department might digitize short (1-2 minutes) video clips of significant findings for use in follow-up examinations to access response to treatment or progression/recurrence of pre-malignant lesions. In the domain of patient education Image Engine could assist healthcare providers in explaining disease processes and procedures using multimedia materials.

Reference

[1]. Huth EJ. The information explosion. Bull N Y Acad Med 1989 Jul-Aug;65 (6):647-61; discussion 662-72

[2]. Huang HK, Aberle DR, Lufkin R et al. Advances in medical imaging. Annals of Internal Medicine 1990 Feb 1;112(3):203-20

[3]. Gur D, Straub WH, Lieberman RH, Gennari RC. Clinicians' access to diagnostic imaging information at an academic center: perceived impact on patient management. American Journal of Roentgenology 92 Apr;158(4):893-6

[4]. Seshadri SB, Arenson RL. The impact of PACS on research and education. International Journal of Bio-Medical Computing 1992 May; 30(3-4):263-6

[5]. McArthur JR, Bolles JR. A survey of videodisc and interactive videodisc projects in North America: Part I. Journal of Audiovisual Media In Medicine 1992 Apr;15(2):61-72

[6]. Ackerman MJ. The videodisc in medical education. Radiographics 1992 Jan;12(1):121-2

[7]. Arenson RL. Picture Archiving and Communication Systems. Western Journal of Medicine 1992 Mar;156(3):298-9

[8]. Schmiedl UP, Rowberg AH Literature review: picture archiving and communication system. Journal of Digital Imaging 1990 Nov;3(4):238-45

[9]. Blair DC, Maron ME. An evaluation of retrieval effectiveness for a full-text document-retrieval system.Commun. ACM 28, 3 (mar. 1985), 298-299

[10]. Lowe, H.J.; Barnett G.O. MicroMeSH: A Microcomputer System for Searching and Exploring the National Library Medicine's Medical Subject Headings (MeSH) Vocabulary. Proceeding of Eleventh Annual Symposium on Computer Applications in Medical Care. IEEE, Washington, D.C. 1987

[11]. Lowe HJ. AURIC: A prototype object-oriented UMLS browser and query generator. Second Annual Educational and Research Conference, American Medical Informatics Association, San Francisco CA, June 1991